System Health Monitoring for Space Mission Operations

David L. Iverson
NASA Ames Research Center
Mail Stop 269-4
Moffett Field, CA 94035
650-604-3115
David.L.Iverson@nasa.gov

Abstract— Many spacecraft provide an abundance of system status telemetry that is monitored in real time by ground personnel and archived to allow for further analysis. In the flight control room, controllers typically monitor these values using text or graphical displays that incorporate individual parameter limit checking or simple trend analysis. Recent developments in data mining techniques for anomaly detection make it possible to use the wealth of archived system data to produce more sophisticated system health monitoring applications. These "data driven" applications are capable of characterizing and monitoring interactions between multiple parameters and can complement existing practice to provide valuable decision support for mission controllers.

Data driven software tools have been successfully applied to mission operations for both the Space Shuttle and the International Space Station. These tools have been applied to engineering analysis of spacecraft data to detect unusual events in the data, and to real time system health monitoring in the flight control room. Augmenting traditional mission control software with advanced monitoring tools can provide controllers with greater insight into the health and performance of the space systems under their watch. Adding heuristic rule based methods that encode system knowledge obtained from seasoned mission controllers can also be helpful to less experienced personnel. We will describe how such techniques have been applied to NASA mission control operations and discuss plans for future mission control system health monitoring software systems. [1]

TABLE OF CONTENTS

1. INTRODUCTION

Most NASA mission controllers are responsible not only for operating their designated spacecraft subsystems to meet mission objectives, but also for monitoring those subsystems to ensure that they are operating properly. Insight into subsystem performance is provided by down linked telemetry data produced by sensors on board the spacecraft. In the flight control room, controllers typically monitor these values in real time using text or graphical displays that incorporate individual parameter limit checking or simple trend analysis. Mission control engineering support groups are available to perform more in depth analysis of collected telemetry to assess system health or to explain unusual phenomena observed in system behavior.

Recent developments in data mining techniques for anomaly detection make it possible to use the wealth of archived spacecraft system data to produce advanced system health monitoring applications that can aid mission controllers and engineering analysts in their tasks. These "data driven" applications are capable of characterizing and monitoring interactions between multiple parameters and can provide valuable decision support for mission controllers and engineers.

Several data driven software tools, including Orca and the Inductive Monitoring System (IMS), have been successfully applied to mission operations for both the Space Shuttle and the International Space Station. Orca [1] is a data mining tool that searches for unusual data points, or outliers, in multivariate data sets by calculating the distance of each data point from neighboring points. The presence of outliers in spacecraft system data is of interest to mission controllers because they may indicate malfunctioning system components. The IMS tool [2] uses a data mining technique called clustering to analyze archived spacecraft data and characterize nominal interactions between selected parameters. This characterization, or model, of normal operation is stored in a knowledge base that can be used for real time system monitoring or analysis of archived events. Spacecraft data is compared with the nominal model built by IMS to produce a measure of how well the data matches the normal behavior defined by the training data. Significant deviations from the nominal system model can alert the controller to a system malfunction or precursor to a significant failure.

Supplementing current mission control software with advanced monitoring tools, such as Orca and IMS, can provide controllers and engineers with greater insight into the health and performance of the space systems under their watch. The addition of heuristic rule based methods that encode system knowledge obtained from seasoned mission controllers can also be helpful to less experienced personnel. We will discuss how these tools and techniques have been used to complement existing NASA mission control practices and present plans for future mission control system health monitoring software deployments using these technologies.

2. Standard MCC Software Tools

The NASA Mission Control Center (MCC) at the Johnson Space Center (JSC) employs a standard suite of certified software tools in the flight control rooms. Individual controller disciplines may also introduce specialized software tools suitable for their particular task. Core software applications used for data monitoring in the MCC include the Information Sharing Protocol (ISP) data distribution system, the MSK-View and RTPLOT data display tools, the ELOG event logging utility, and the ISPATOM computation tool.

Information Sharing Protocol (ISP)

The Information Sharing Protocol (ISP) software is a distributed system that supports real-time telemetry distribution and messaging in the MCC. [3] It is the primary spacecraft data feed for the mission control consoles. In addition to subscribing to telemetry data, console applications can publish values to ISP to allow other software applications to subscribe to and use locally computed results along with the real-time telemetry data.

MSK-View

MSK-View is a program used to display ISP data in a tabular format defined in a screen configuration file. [4] Figure 1 shows an example of a display created with MSK-View. It allows the user to define screens with alphanumeric labels and parameter values along with simple line drawings and screen navigation buttons. Each object on the screen can be assigned a color as well. The MSK-View application can subscribe to an ISP data feed and provide a text display of real-time data values as they are received. In addition, the display can show different user defined text strings based on the value of a data parameter. For instance, a display element can be defined to display "HIGH" if a data value exceeds a high limit value, and "LOW" if it is less than a low limit value. This gives controllers a way to detect off nominal parameters on their MSK-View screens using predetermined limit values.

RTPLOT

The RTPLOT (Real Time Plot) program provides a means of plotting real-time or played-backed data accessed via the ISP data feed. [5] RTPLOT subscribes to the ISP data stream and provides line plots of one or more parameter values over time. (Figure 2). It allows the user to define upper and lower limits for each plot, which appear as red dashed lines on the plot display. The plot can be iconified to a smaller size to conserve screen real estate. If a parameter on an iconified plot falls outside a limit, the plot name on the icon will be displayed in red to alert controllers of the out of limit condition. RTPLOT also includes routines to estimate future data values by fitting lines or parabolic curves to previously plotted data.

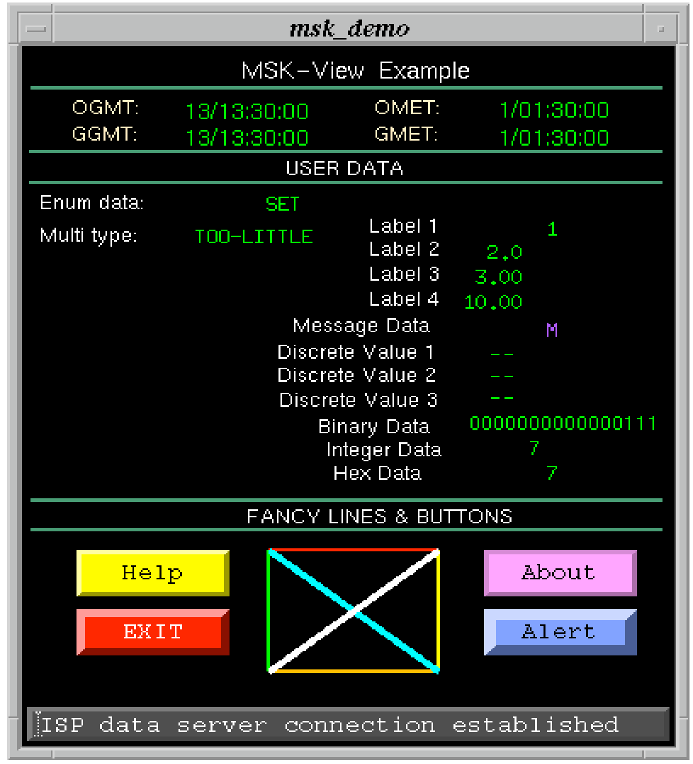

Figure 1: Sample MSK-View Screen

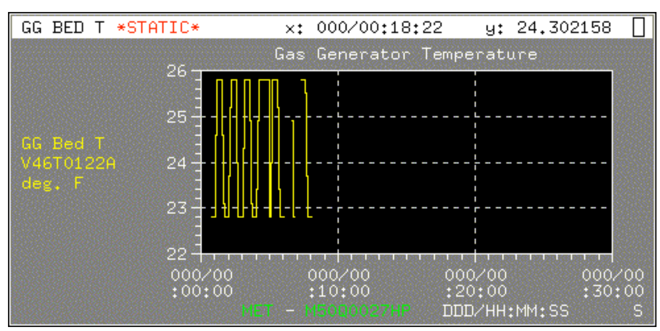

Figure 2: Sample RTPLOT Display

ELOG

The ELOG event logger program provides a means to automatically log telemetry change events in real time. [6] ELOG defines an event as a comparison (greater than, equal to, etc.) between an ISP parameter value and a constant. When the comparison is true for a specified amount of time, the event is logged to a file and displayed in a list on the controller's screen. The display will show the time the event occurred and a text message describing the event with a user selected background color. (Figure 3) The controller can acknowledge that they have seen an event in the ELOG display by clicking on that list item with the mouse. An acknowledged event will be displayed in reverse video. Audible alarms can be associated with critical events to attract the controller's attention. Like RTPLOT, the ELOG display windows can be iconified to save screen space. An ELOG icon will show the name of the window, a count of how many events have been logged in that window, and a count of how many of those events have not been acknowledged.

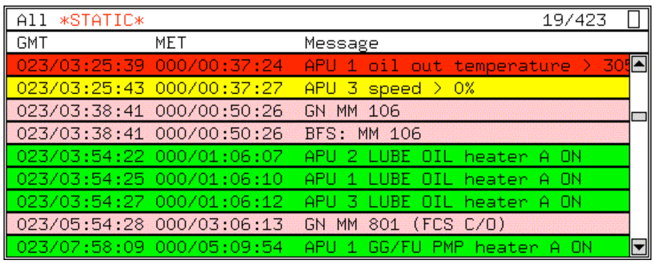

Figure 3: Sample ELOG Display

ATOM

The Advanced Tool of Math (ATOM) is a software tool that allows controllers to define and execute real time computations on telemetry data by specifying equations using ISP parameter identifiers. [7] ATOM will parse the user defined equation, subscribe to the required real-time telemetry values, calculate results as new data arrives, and publish the results back to ISP for use by other applications. ATOM provides common mathematical and comparison functions, including square root, logarithmic, and trigonometric functions. It can also perform calculations on time series data, such as average, standard deviation, and future value prediction using a least squares line fit. Program flow control is provided by IF, WHILE, and CASE statements. Basic data structures like arrays and stacks can also be used in the calculations.

3. Data Mining for Mission Control

NASA maintains years of archived Space Shuttle and International Space Station (ISS) telemetry data in the Operational Data Reduction Complex (ODRC) at the Johnson Space Center. This archive can provide a wealth of information about the behavior of individual telemetry parameters and how those parameters correlate to each other. The standard suite of mission control software tools is able to convey current data and system status to controllers, but takes little advantage of information available on the historic behavior and interactions of the parameters.

Methods from the field of data mining are useful for analyzing and characterizing the type of data found in the ODRC archive. In particular, recent developments in data driven anomaly detection techniques can process the data to find unusual events, or outliers, in data for a given subsystem. These anomaly detection techniques can also automatically analyze archived nominal system data to characterize normal system performance. Comparing incoming real-time data to that nominal model can let the user know if the current system behavior differs from previous system performance.

One powerful feature of these data mining techniques is the ability to simultaneously analyze multiple parameters. This allows them to discover and model interactions between related parameters that might be difficult to notice when monitoring those parameters individually.

4. ORCA Distance-Based Outlier Detection

Orca is a data mining tool that analyzes multivariate data sets. [1] It uses a nearest neighbor approach to outlier detection. Each group of parameter values is considered a vector that defines a point in a multi-dimensional vector space. For each point in the data set, Orca locates other points in the data set that are closest to that point, called the point's nearest neighbors. Distance between points is measured with the Euclidean distance measure. The program outputs a score for each point representing the average distance to the nearest k neighbors in the data set. The value of k is specified by the user. Points that have a larger average distance to their nearest neighbors than most other points in the data set are considered outliers. The Orca program is able to find outliers within a single data set, or compare one data set with another to determine which points in the first data set are unusual in comparison to the second set.

One approach to using Orca with spacecraft telemetry is to form data vectors by grouping a number of concurrent parameter values collected from sensors on a given spacecraft subsystem and searching for outliers among those vectors. For instance, vectors could be formed from temperature, pressure, and fuel flow rates in a rocket engine. Values collected simultaneously at regular time intervals from each of the sensors would form the data set. An Orca analysis of this data set can locate data from time periods during the engine firing that display unusual characteristics compared with the rest of the data. Those anomalous data points may be symptoms of engine malfunctions, such as a faulty pressure regulator or an incorrect fuel-oxidizer mixture ratio.

5. IMS: Inductive Monitoring System

The Inductive Monitoring System (IMS) is a tool that uses a data mining technique called clustering to extract models of normal system operation from archived data. [2] Like Orca, IMS works with vectors of data values. IMS analyzes data collected during periods of normal system operation to build a system model. It characterizes how the parameters relate to one another during normal operation by finding areas in the vector space where nominal data tends to fall.

These areas are called nominal operating regions and correspond to clusters of similar points found by the IMS clustering algorithm. These nominal operating regions are stored in a knowledge base that IMS uses for real-time telemetry monitoring.

During system monitoring, IMS reads real-time or archived data values, formats them into the predefined vector structure, and searches the knowledge base of nominal operating regions to see how well the new data fits the nominal system characterization. For each input vector, IMS returns the distance that vector falls from the nearest nominal operating region. Data that matches the normal training data well will have a deviation distance of zero. If one or more of the data parameters is slightly outside of expected values, a small non-zero result is returned. As incoming data deviates further from the normal system data, indicating a possible malfunction, IMS will return a higher deviation value to alert users of the anomaly. IMS also calculates the contribution of each individual parameter to the overall deviation, which can help isolate the cause of the anomaly.

6. Mission Control Applications

The Orca and IMS tools have both been applied in NASA mission control to support real-time telemetry monitoring and engineering analysis of mission data. In support of the JSC Mission Evaluation Room (MER) engineering analysis activity, the tools were applied to data from the Space Shuttle Wing Leading Edge Impact Detection System (WLEIDS) to find potential impact signatures. In the International Space Station (ISS) flight control room they have been used to build real-time health monitoring applications for the ISS Control Moment Gyroscopes.

Space Shuttle Wing Leading Edge Impact Detection System

The Shuttle WLEID system was developed in response to the loss of the Columbia orbiter on the STS-107 mission. During the launch of STS-107 a piece of foam shed from the Shuttle external fuel tank struck the leading edge of the orbiter's left wing, compromising the thermal protection system. This damage resulted in the tragic loss of vehicle and crew during reentry due to overheating and failure of the internal wing structure. [8]

The WLEIDS consists of 132 single axis accelerometers mounted along the length the orbiter leading edge wing spars. (Figure 4) During launch, the accelerometers collect data at a rate of 20 KHz and store that data onboard for subsequent downlink to Mission Control. Within 6 to 8 hours of launch, summary files containing periodic sub-samples of the data collected by each accelerometer are down linked to the MER for analysis to find potential impact signatures. This analysis must be performed within 24 to 48 hours of the launch so the results can be used to schedule

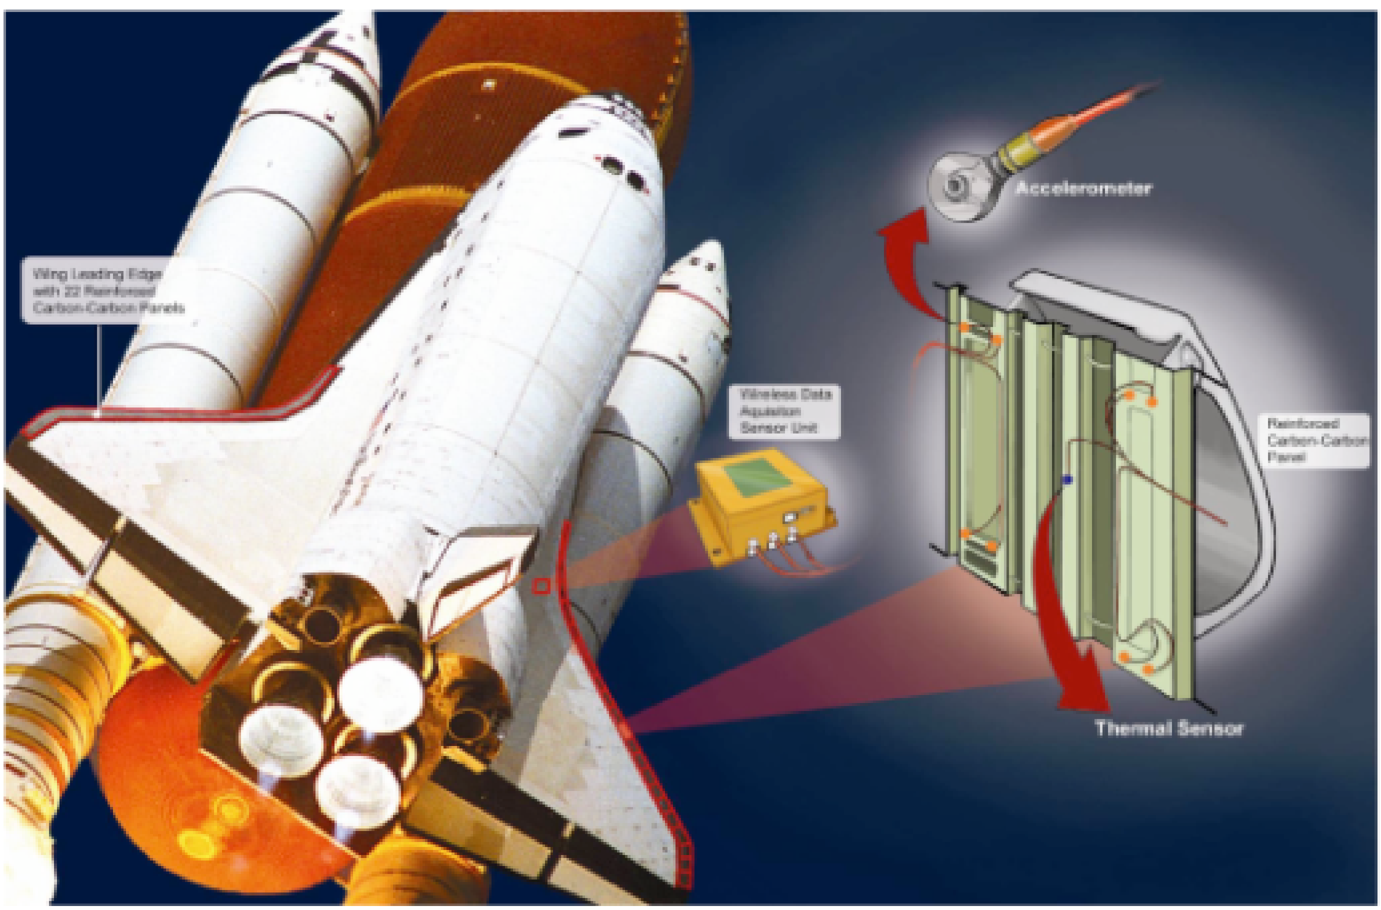

Figure 4: Space Shuttle Wing Leading Edge Impact Detection System

detailed on-orbit wing leading edge inspections using cameras mounted on the Shuttle robotic arm.

The basic WLEIDS analysis is performed by MER engineers by visually examining three dimensional graphs of the summary data that show accelerometer location and vibration magnitude along a time axis. (Figure 5). The analysts search the graphs for localized peaks among the normal vibration signals caused by the Shuttle engines and aerodynamic forces. Unusual peaks in the data might have been caused by an impact on the wing leading edge. When possible impact events are identified, a half second of raw data collected by the affected accelerometer at that time is downloaded for more thorough analysis.

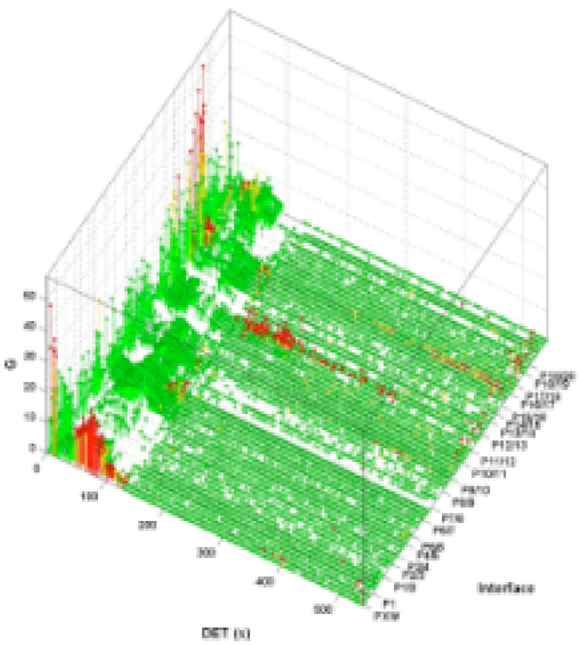

Figure 5: Sample WLEIDS Summary File Graph

The Orca and IMS tools have been used to support the WLEIDS analysis on three Shuttle launches. The goal is to provide a quick initial scan of the WLEIDS summary files to locate unusual points and help focus the MER analysts' efforts. For each accelerometer sensor, vectors were formed from concurrent values collected from that sensor and seven nearby sensors. Orca is used to search for outliers within the data collected during the current launch. Prior to the launch, IMS was used to analyze normal data from previous launches to characterize typical vibration patterns for each group of accelerometers. Data from the current launch is compared to this characterization to identify unusual vibration patterns that might have been caused by impact events.

To gauge the effectiveness of the data mining tools, we compared Orca and IMS results to visual WLEIDS summary file analysis performed by MER engineers on the STS-115 launch of Space Shuttle Atlantis. The analysts classify the events they identify as critical, probable, or questionable based on the likelihood that the data signature was caused by an impact. Their analysis of the STS-115 launch WLEIDS summary data produced 6 critical events, 23 probable events, and 2 questionable events. The Orca analysis placed all critical events in the top 50 outliers. IMS identified 334 interesting events, divided nearly evenly between the two wings. Those events included all 6 critical events, 18 of 23 probable events, and all of the questionable events found by the MER analysts. Most of the anomalies identified by Orca and IMS that were not noted by the analysts could be eliminated as normal global vibrations that shook the entire vehicle. Additionally, during all launches where Orca and IMS have been used, the tools identified several lower energy vibration signatures that did not stand out in the visual data inspection. These events were investigated with raw WLEIDS data downloads from the affected sensors. Fortunately all of the potential impact events identified in the WLEIDS summary data were shown to be the result of non-damaging phenomena, such as aerodynamic events or sensor data spikes, and all missions concluded with safe and uneventful reentry and landing.

ISS Control Moment Gyroscopes

The International Space Station (ISS) Control Moment Gyroscope (CMG) attitude control system consists of four large gyroscopes, each mounted in a gimbal system that can rotate the CMG about the two axes perpendicular to the gyroscope spin axis. (Figure 6) The CMGs operate as non-propulsive attitude control devices that exchange momentum with the ISS through induced gyroscopic torques.

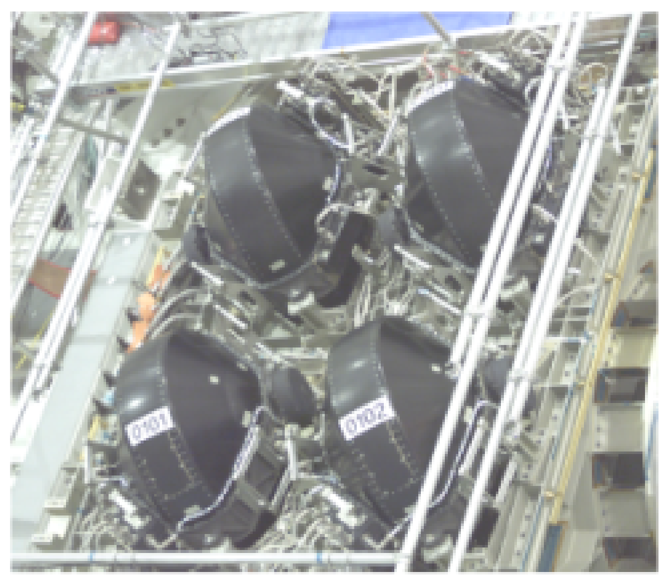

Figure 6: ISS Control Moment Gyroscopes

As they have aged some of the CMGs have degraded enough to malfunction and require replacement. A failed CMG1 was replaced with a new unit in July 2005, and a faulty CMG3 was replaced in August 2007. Given their history, the ISS Attitude Determination and Control Officer (ADCO) flight controllers are interested in detecting early symptoms of degradation in the CMGs. A deployment of data driven system health monitoring applications in the ISS flight control room is assisting with that task.

Working with the ADCO flight controllers, 13 CMG parameters were selected for real time monitoring. These parameters include CMG vibration, bearing temperatures, rotation speed, gimbal rates, electrical current, and ISS rotation rates, along with rates of change of the temperatures and electrical current. Archived data collected over a period of 10 months for CMG1, 2, and 4 was analyzed. Six weeks of available data for the recently installed CMG3 was analyzed for that unit. The data was sampled at a 1 Hz rate and formed into vectors of 13 values. Each CMG was analyzed individually to capture its unique characteristics.

The first operation with the CMG data was to remove any anomalies from the archived data. This was accomplished by searching for outliers within each data set using the Orca tool. Data records with significant deviations relative to the remainder of the data for that CMG were removed. These deviations were typically caused by data corruption or minor anomalies in CMG operation. Once the archived CMG data had gone through this cleaning process, the remaining nominal data was used by IMS to build a monitoring knowledge base for each CMG.

The IMS monitoring application was integrated with the ISP data server software to access real-time telemetry in the ISS flight control room. Four IMS processes, one per CMG, are run on the ADCO flight control console to provide continuous monitoring. Once per second, when data is available, each IMS process will query the appropriate CMG knowledge base and return the amount of overall deviation, if any, from the nominal training data. It will also return the contribution of each individual parameter to any deviation to aid in isolating the source of the deviation. These results are published back to the ISP data stream for access by other software applications.

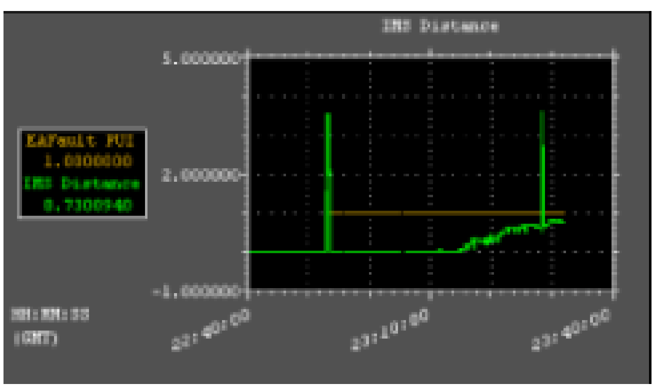

Figure 7: IMS Detection of a CMG Electronics Fault

The ADCO controllers have developed RTPLOT displays to graph the IMS CMG results in real-time on their control console. These graphs read the IMS results from the ISP server and show the amount of total IMS deviation over time for each CMG (Figure 7), along with the individual parameter contributions. The controllers also use ELOG rules to notify them when IMS detects unusual CMG behavior. The IMS output acts as a decision aid, alerting the controllers to CMG anomalies and prompting them to investigate unusual occurrences. It also serves as a cross check for atypical events detected by other means. For example, when a controller sees a pattern on their telemetry display that is different from what they expect, they can check the IMS output for that time period to see if this type of behavior had occurred previously. If IMS shows minimal deviation during that time, then the current data matches previous behavior seen in the nominal IMS training data. If IMS shows a significant deviation during that time, further investigation may be warranted.

7. Heuristic Fault Identification

The IMS processes monitoring the ISS Control Moment Gyroscopes can alert the controllers when a CMG deviates from typical behavior, but the basic IMS routines do not identify the fault that caused the deviation. The ADCO controllers would like to receive explicit alerts in the event of particular CMG faults when possible. Some of these fault signatures are straightforward enough to encode as monitoring rules and integrate with the IMS console software.

One relatively common, but recoverable, CMG fault is a transient condition in the CMG electronics assembly that the ADCO controllers call an “EA Fault”. This event is characterized by a brief electrical current spike beyond a threshold value accompanied by the setting of an error bit in the electronics assembly. A routine was added to the IMS monitoring software that watches for the occurrence of an appropriately sized current spike and the setting of the electronics assembly fault bit within a short time window. If both of these events occur within that time span, IMS will set an EA Fault bit on the ISP data stream to alert the controllers of the incident.

Another less frequent incident is a Loss of Comm fault that occurs when a CMG experiences problems with the network connection to the ISS computer responsible for commanding the unit and gathering CMG sensor data. This event is automatically recognized by ISS on board fault detection logic, which initiates a recovery procedure to reestablish the connection. This recovery procedure results in transmission of several specific status messages to the ground controllers. Watching for these status changes allows the IMS Loss of Comm detection routine to detect when the event has occurred and set the appropriate alert bit in the ISP data server. The routine also detects when a successful recovery has occurred and resets the alert bit.

Although these heuristic fault identification abilities do not cover all possible fault scenarios, they add value by enabling the controller to quickly determine the cause of these more common anomalies. This can be particularly useful for less experienced controllers that may not have seen these particular fault signatures. Additional fault detection routines may be incorporated in the future.

8. Summary and Future Work

Through practical application, it has been demonstrated that data driven system health monitoring applications can be useful in a space mission operations setting. Many spacecraft have extensive archives of telemetry data available that can be advantageously exploited by data mining methods. Two data mining tools, Orca and the Inductive Monitoring System, have been used to analyze data from the Space Shuttle and International Space Station to search for anomalous data points that could be indications of a fault or damage to the spacecraft. Providing information on possible system anomalies in a timely manner provides controllers and mission support engineers with helpful decision support and enables more efficient and effective execution of their duties. The inclusion of fault identification routines further simplifies their tasks.

These initial tool deployments in the NASA JSC mission control center have demonstrated the utility and effectiveness of data driven system health monitoring methods in two disciplines, but the applications are not limited to just these examples. There are many areas with rich archived data repositories where these and similar techniques can be applied. Mission controllers from several ISS disciplines, including power management, thermal control, and life support, have expressed interest in similar system monitoring tools. Now that the software has been integrated with the mission control data systems, expansion of the capability is primarily a matter of identifying relevant parameters to monitor and performing the archived data analysis. Work is beginning now to develop monitoring capability for ISS thermal control systems.

Eventually we plan to develop tools that allow mission controllers to build their own data driven monitoring applications. They will specify which parameters to monitor, what time periods to include in the nominal training data, and any computations that should be performed on the raw telemetry data. The tool set will retrieve the desired archived training data, remove spurious data points using outlier detection, and build a new monitoring knowledge base and an appropriate monitoring application configuration to run on their control console.

A useful enhancement to the current monitoring software would be the ability to automatically detect operating mode changes in the monitored system and switch to a targeted monitoring knowledge base developed specifically for that mode. For instance, the ISS is flown in different orientations and configurations during different mission phases. The behavior of the CMGs can differ in the various configurations. Rather than building one large knowledge base per CMG that covers all cases, as in the current deployment, a separate knowledge base could be built from archived data collected during each ISS configuration, then consulted for real-time monitoring when the ISS is in that configuration. This would provide more accurate and efficient monitoring capability.

Another application of data driven monitoring to explore is the use of supervised learning methods to help identify fault signatures. If examples of fault behavior are available in the archived data, supervised learning algorithms, such as decision tree or support vector machine based techniques, may be able to analyze the data and distinguish between different types of fault behavior and normal operation. If the monitored system exhibits unusual behavior, fault characterizations from the supervised learning algorithm could help controllers identify the cause of the anomaly. These techniques could allow automated fault identification in cases that are too complex to be encoded in heuristic rules.

Acknowledgements

The author would like to thank the following collaborators who directly supported and contributed to the development and deployment of the mission control software described in this paper: From NASA Ames Research Center – Rick Alena, Charles Lee (EASI), Dawn McIntosh, Mark Schwabacher, Lilly Spirkovska, and William Taylor. From NASA Johnson Space Center – Tatiana Dobrinskaya (USA), Ann Esbeck (USA), Tammy Gafka, Joshua Johnson, Keith Martin (TSI), and Charles McCann.

References

[1] Stephen D. Bay and Mark Schwabacher, "Mining Distance-Based Outliers in Near Linear Time with Randomization and a Simple Pruning Rule," *Proceedings of The Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,* 2003.

[2] David L. Iverson, "Inductive System Health Monitoring," *Proceedings of The 2004 International Conference on Artificial Intelligence (IC-AI04)*, Las Vegas, Nevada, June 2004. CSREA Press.

[3] Mission Control Center Platform Services System Information Sharing Protocol User's Guide, NASA Document CSOC-JSC-UG-002098, April, 2002.

[4] MSK-View Users Guide, Version 3.4, September 1997.

[5] RTPLOT Reference Guide, NASA Operations Document JSC-26101, August 1998.

[6] ELOG Reference Guide, NASA Operations Document JSC-26101, July 1998.

[7] Adam Dershowitz, Advanced Tool of Math (ATOM) Users Guide, Version 02.03.06, July 2006.

[8] Columbia Accident Investigation Board Report, Volume 1, August 2003.

BIOGRAPHY

David Iverson *has worked as a computer engineer in the Intelligent Systems Division at the NASA Ames Research Center for more than two decades. During most of that time, he has been involved in integrated system health management (ISHM) for spacecraft, aircraft, and scientific payload systems. He has investigated several aspects of this field including reasoning from system fault models, model based monitoring and diagnosis, and data driven techniques for system monitoring. Most recently, he has developed and applied the Inductive Monitoring System data driven anomaly detection software. He is the NASA Ames project lead for deployment of data driven monitoring for JSC mission control applications. He received his B.S. from the University of Puget Sound where he studied Computer Science, Mathematics, and Physics. He also holds an M.S. in Computer Science from Stanford University.*